PA-8500 to Hit 440 MHz

Although the PA-8500 has slipped into early next year, HP has countered by boosting its clock speed. The company will initially deliver 360-MHz parts but expects to quickly boost the clock speed to 440 MHz. The speedup comes from a slightly modified design with some critical speed paths improved. Both will ship in January; HP had previously committed to shipments in 1998.

The increase in clock speed is fortuitous, as the initial part will fall far short of the target of 30 SPECint95 (base) and 50 SPECfp95 (base) that HP set for itself last year (see MPR 12/29/97, p. 1). The 440-MHz part, however, will deliver an estimated 30 int/50 fp, exactly on target.

These scores are greater than any yet reported, passing the Alpha 21264 (see MPR 9/14/98, p. 4) at 26 int/41 fp. The Alpha chip is now slated for December shipments, giving Compaq little time to boost the performance of its systems before the 440-MHz PA-8500 begins shipping.

In a bid to capture MPR's 1998 award for Highest Manufacturing Cost, the corpulent PA-8500 weighs in at 475 mm^2 in 0.25-micron technology—a surprise, given HP's prior statement that the PA-8500 would be only "somewhat larger" than the 345-mm^2 PA-8200. Even with a generous yield credit for SRAM redundancy, we estimate the manufacturing cost to be on the far side of $400. While this cost is much higher than that of other processors, the PA-8500 requires no expensive cache SRAMs, as it includes 1.5M of on-die L2 cache. HP did not disclose its fab; Intel is a possibility.

The 360-MHz PA-8500 will begin shipping in January as a board upgrade to HP's Visualize C200 and C240 workstations. The PA-8500 board lists for $8,000—steep, but a far cry from the $65,000 for an Alpha 21264 upgrade board. HP is likely to offer the faster PA-8500 in new workstation and server products to be announced later this year. The new PA processor should keep HP dueling with Alpha in the performance race, at least until Merced appears. —*L.G.*

Power3 Meets Targets, Aims for 500 MHz

In a clear violation of the microprocessor version of the Pauli exclusion principle—which states that no new microprocessor can achieve its predicted schedule, die size, *and* performance—IBM's Power3 has done just that. The principle has gained credence of late, as Intel pushed Merced out a year and down in performance expectations, the Alpha 21264 missed its FP performance goal by 25%, HP's PA-8500 came in 35% too big, and UltraSparc-3 and the MIPS R12000 were both delayed by at least nine months.

But at last year's Microprocessor Forum, IBM said that Power3 would be available in systems in 2H98 at 200 MHz with performance of 12 SPECint95 (base) and 28 SPECfp95 (base). Sure enough, the company has made good on that promise by announcing availability of both a single- and dual-processor version of its Model 260 workstation based on the 200-MHz Power3 chips.

Admittedly, IBM did not set the bar high. At 200 MHz, Power3's 12.5 SPECint95 (base) score puts it below even low-cost PCs. Floating-point performance is more respectable, delivering 640 MFLOPS on 1,000 × 1,000 Linpack benchmarks and 27.6 on SPECfp95 (base)—but still well behind the 21264's score of 41 and the PA-8500's estimate of 50. Of course, Power3 is shipping and the others are not.

Memory bandwidth is also more competitive: with a peak memory bandwidth of 1.6 Gbytes/s and an L2 cache bandwidth of 6.4 Gbytes/s, Power3 can sustain 1.1 Gbytes/s on the Stream benchmark. The part's strong floating-point performance and outstanding memory bandwidth give it industry-leading 3D performance, scoring 218 on the CDRS (OpenGL) benchmark.

At the Microprocessor Forum earlier this month, IBM demonstrated a new 400-MHz version of the chip that it calls Power3-2. The 100% boost in frequency was achieved by tuning critical speed paths and rendering the part in the company's 0.22-micron copper CMOS-7S process. The newer version improves on the original by adding four-way set-associativity to the L2 (previously direct mapped), which should improve its performance on commercial database applications. New fractional bus dividers allow CPU frequency to be increased in smaller steps without sacrificing memory bandwidth. IBM expects to deliver workstations based on Power3-2 in 2H99.

In 2000, IBM plans to bump the frequency to 500 MHz by employing the company's 0.22-micron silicon-on-insulator technology (see MPR 8/24/98, p. 8). But this will be the last in the line of the Power3 processors. Quietly, IBM admits the error of sacrificing CPU frequency to achieve greater instruction-level parallelism.

The company is now developing a completely new microarchitecture that avoids this pitfall. Code-named Giga Processor, the design targets operation at over 1 GHz and is rumored to include two CPUs on one chip. The part aims to compete with Intel's Merced in late 2000 for the high-end server market. —*K.D.*

Intel Ships 450-MHz Xeon for Workstations

Making a bigger push for its Pentium II Xeon in workstations, Intel has increased the speed of the entry-level product in that line and cut the price. The company has essentially replaced the Xeon-400/512K with a Xeon-450/512K, introducing the new product at a list price of $824 (in 1,000s), well below the $1,124 price it had been asking for the previous part. Intel hopes the lower price will encourage more workstation vendors to use the Xeon part.

Xeon's biggest competition is Intel's own unenhanced Pentium II. The standard Pentium II-450 began shipping a

month before the Xeon version and is currently priced at $562. Thus, even with the Xeon price cut, Intel is still charging a 47% premium for nothing more than running the L2 cache at full speed rather than half speed. The faster cache has little impact on workstation applications: 4% on SPECint95 and 9% on SPECfp95, for example. But workstation makers wanting the fastest x86 processor available are willing to pay a few hundred dollars more to get even a small edge.

Intel plans to deploy Xeon-450 processors with 1M and 2M caches, but not until early next year. The company blames the delay on testing for four-CPU configurations, which it doesn't do on the 512K parts. In the interim, Intel cut the price of the Xeon-400/1M to $1,980 from $2,836.

We do not expect Intel to institute regular quarterly price cuts for the Xeon line as it has for its PC processors. Instead, the company is likely to bring in future faster parts at the same price as the current parts, much as it did with the Xeon-450/512K. This month's changes look to be a one-time adjustment to increase the popularity of the Xeon parts by reducing the initial outrageously high prices. —*L.G.*

Intel Converts to 100% 0.25 Micron

Intel's latest round of price cuts obsoletes all of the older 0.28-micron products and shows the company rapidly driving its Pentium II line to a 100-MHz bus from top to bottom. The least expensive Pentium II is now the 333-MHz Deschutes at $181. The 300-MHz Klamath remains on the price list, but at a noncompetitive price of $192; all other Klamath parts are gone. Intel confirmed that production of the 0.28-micron Klamath has stopped.

The 350-MHz Pentium II is now only $213, 50% less than it cost just one quarter ago. Intel expects most PC makers to ignore the 333-MHz Pentium II in favor of either the 350-MHz version, which costs only $32 more, or the Celeron-333, which costs $22 less and has nearly the same performance. Intel also dropped the Covington-266 from its price list and is dumping the 300-MHz version at $95 (see MPR 10/26/98, p. 39 for details). Mobile prices remain at the same levels announced last month (see MPR 9/14/98, p. 4). —*L.G.*

iCompression Offers Live(ly) MPEG-2 Encoder

Startup iCompression (*www.icompression.com*) has provided the first technical details of its Vivace-izC, the world's first single-chip real-time MPEG-2 video and audio encoder to also handle "system-level" encoding—the coding of the transport or program streams. With other encoders, such as the DVx family from C-Cube Microsystems (see MPR 12/8/97, p. 1), audio and system-level encoding must be handled by separate chips.

Speaking at the Microprocessor Forum earlier this month, iCompression founder Govind Kizhepat described how the new chip performs essentially all the functions needed for real-time DVD encoding from any source of standard video and audio. The chip accomplishes this goal with just 10M of external SDRAM (an 8M bank for the video engine and 2M more for audio and system data) plus a low-cost microcontroller to handle user-interface functions.

The Vivace-izC includes a pair of DSPs designed by iCompression to execute its internally developed video and audio compression algorithms. The DSPs execute Java bytecodes plus iCompression's own multimedia extensions. They are assisted by hardwired DCT (discrete cosine transform) and inverse DCT engines, a motion-estimation unit, digital video filters, and other fixed-function logic.

The chip comprises 6.1 million transistors (65% logic, 35% SRAM) implemented in 0.35-micron technology and operating at 90 MHz on a 3.3-volt supply. The Vivace-izC, packaged in a 420-contact BGA, is currently shipping. A PCI-bus evaluation board is also available. Pricing has not been announced. —*P.N.G.*

Virtual Channel SDRAM Supported, Attacked

NEC's Virtual Channel Memory (VCM; see MPR 12/29/97, p. 5), though late to the fray, has gained strategic endorsements from chip-set vendors Acer Labs, SiS, and VIA. At virtually the same time, VCM became the target of a lawsuit filed by Enhanced Memory Systems, supplier of the rival Enhanced SDRAM.

Acer Labs (*www.ali.com.tw*), SiS (*www.sis.com.tw*), and Via (*www.via.com.tw*) will soon offer chip sets that support VC-SDRAM for both Socket 7 and Slot 1 systems. The three supporters of the JEDEC-approved Virtual Channel proposal control almost all of the non-Intel PC core logic market.

The new chip sets will use 133-MHz Virtual Channel SDRAM to achieve 50% better effective bandwidth than current 100-MHz PC100 SDRAM. Part of the improvement comes from the faster clock rate; the rest comes from the Virtual Channel protocol and chip-design enhancements that improve bus utilization and reduce page-miss rates.

NEC is currently sampling 64-Mbit VC-SDRAMs in 4-, 8-, and 16-bit bus widths, with production scheduled for this month. NEC has signed up Siemens as a second source for the new DRAMs, and it is working to recruit at least two more alternate sources.

The new chip sets are the first products with 133-MHz processor-bus speeds, matching the performance of the new VC-SDRAM devices. The faster bus, however, may be moot: although Intel will deploy a 133-MHz Slot 1 next year, no vendor has announced plans for 133-MHz Socket 7 processors. All of the VC-SDRAM chip sets will include AGP support, though some will have integrated 3D graphics (see MPR 8/24/98, p. 4) in lieu of external AGP interfaces.

In a separate announcement that same day, Enhanced Memory Systems (*www.edram.com*) claimed that NEC's VC-SDRAMs uses technology covered by U.S. patent number 5,721,862, filed in 1995 by Enhanced's parent company,

Ramtron International. The patent covers a variety of DRAM architectures with row-data latches or caches. Enhanced seeks an injunction and an unspecified amount of damages.

The Enhanced SDRAM (ESDRAM), manufactured by IBM and Siemens, is also a JEDEC-approved superset of the SDRAM. Though ESDRAM is supported by various embedded-processor memory controllers as well as VIA's Apollo Pro and MVP-3 chip sets, Enhanced has yet to receive an endorsement comparable to the Virtual Channel announcement.

Though this skirmish is significant, it is at best a battle for second place. ESDRAM and VC-SDRAM each have just two sources today. Acer Labs, SiS, and VIA together represent about a third of the market for PC core logic, and systems based on their chip sets—typically aimed at the low end—represent an even smaller portion of the market for main memory. Intel, with its majority of the PC core-logic business, has endorsed the Rambus Direct RDRAM, which will be offered by all major memory vendors.

Hopefully, Enhanced, NEC, and the non-Intel chip-set vendors (and those behind the SynchLink DRAM effort) can resolve this conflict. One mutually acceptable alternative to Direct RDRAM may succeed; three competing alternatives will not. As Ben Franklin might have said, they must hang together, or they will certainly hang separately. —*P.N.G.*

SGI, Real3D Swap Suits for Ties

A new strategic relationship between SGI and Real3D settles three years of pending litigation and gives SGI a stake in the Lockheed Martin subsidiary. SGI purchased "a little bit less than 10%" of Real3D, according to Gerry Stanley, president of Real3D, which is rumored to be worth about $200 million. SGI thus becomes the third co-owner of Real3D. Intel's investment in Real3D gave it about 20% of the company (see MPR 1/26/98, p. 4); the balance is owned by Lockheed Martin.

The agreement marks the end of a protracted legal battle over 3D rendering technology and gives each company a royalty-free license to the other's patent portfolio.

SGI and Real3D are both strong players in the professional visual-simulation market (for military flight simulation and other applications). While the two companies have not announced any plans to codevelop or cross-market their professional 3D products, it is likely that such opportunities are being discussed.

Real3D's experience with mainstream PC graphics, through its role in the development of Intel's 740 graphics chip, is only indirectly relevant to SGI's plans in the Windows NT space, but the two companies may find ways to work together in these markets as well —*P.N.G.* M